

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In Application of : **AFEK et al**

Serial No.: 09/929,877

Group Art Unit: 2151

Filed: August 14, 2001

Examiner: Frantz B. Jean

For: METHODS AND APPARATUS FOR PROTECTING
AGAINST OVERLOAD CONDITIONS ON NODES
OF A DISTRIBUTED NETWORK

APPEAL BRIEF

Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

I hereby certify that this document is being
submitted via EFS Web to the Commissioner for
Patents, P.O. Box 1450, Alexandria, Virginia on
the date shown below.

Dated: 2/11/08 Signature: /David J. Powsner/
David J. Powsner
Reg. No. 31,868

Appellant hereby files its Appeal Brief pursuant to 37 CFR 41.37, a Notice of Appeal having been filed on December 11, 2007. The items required under 37 CFR 41.37(c)(1) are provided in the sections that follow.

This reinstates an appeal that was originally initiated on November 13, 2006, prior to reopening of prosecution. Per MPEP 1204.01, Appellant requests that the Office apply the fees paid under 37 CFR 41.20(b)(2) in the prior appeal (specifically, in connection with the prior Appeal Brief filed January 14, 2007) to this Appeal Brief. Any difference in fees currently required under 37 CFR 41.20(b)(2) and those previously paid are submitted herewith.

(1) Real Party in Interest

The subject application is owned by Cisco Technology, Inc., having a place of business at 170 West Tasman Drive, San Jose, California 95134-1706. The assignment was recorded in the U.S.P.T.O. on February 15, 2006, under Reel 017164, Frame 0886.

(2) Related Appeals and Interferences

Appellant filed a Notice of Appeal in this case on November 13, 2006 and an Appeal Brief on January 14, 2007. That Appeal never reached the Board because the Examiner issued a Notice of Allowability, which was subsequently withdrawn. Appellant is now reinstating the original Appeal pursuant to MPEP 1204.1.

(3) Status of Claims

This application contains claims 1-8, 10, 11, 13-16, 20, 33, 35 and 46-69. Claims 1-8, 10, 11, 13-16, 20, 33, 35 and 46-54 and 56-69 were rejected in an Official Action dated September 11, 2007 (hereinafter referred to simply as the Official Action). Dependent claim 55 was amended on December 11, 2007, to place it in independent form and has since been allowed; *see*, the Advisory Action dated January 24, 2008.

Appellant appeals from the rejection of claims 1-8, 10, 11, 13-16, 20, 33, 35 and 46-54 and 56-69.

(4) Status of Amendments

Appellant filed an Amendment on December 11, 2007, placing claim 55 in independent form and striking paragraphs of the Specification objected to by the Examiner in the Office Action issued on September 11, 2007. Claim 55 has been allowed. Appellants understand, accordingly, that the amendment to the Specification has been entered and that the Examiner's objection to the Specification on the ground that it contained "claims-like language" has been withdrawn.

To the extent that Examiner maintains this objection, Appellants will address it at the conclusion of this appeal. However, if the Appeal Board takes up this issue, notwithstanding that it is merely a formal matter, Appellants submit that (i) the Examiner's objection should be overturned, because there is nothing wrong with submitting "claims-like language" in the disclosure and no basis for objecting to it, and (ii) even if there were proper basis for such an objection, the Amendment submitted by Appellants on December 11, 2007, clearly overcomes the objection by deleting the language in question.

(5) Summary of Claimed Subject Matter

Appellant's invention, as recited in independent claims 1, 46, 49, 56 and 66, provides methods, network elements and systems for responding to and protecting against an overload condition on a network.

Claim 1 recites a method of responding to an overload condition at a "victim" network element in a set of one or more potential victims on a network. The method includes the following steps:

(A) A first set of one or more network elements external to the set of victims is used to initiate diversion of traffic destined for the victim, in response to an indication of an anomalous traffic condition. The network elements in the first set divert the traffic to a second set of one or more network elements external to the set of potential victims. This sort of operation is described (with reference to Figure 1) in paragraph 0248, for example: Routers R0-R8 selectively divert traffic destined for a victim H0 to guards G0-G3 when the victim comes under an anomalous traffic condition. (Paragraph numbers refer to the published version of the present patent application, US 2002/0083175.)

(B) The elements in the second set filter the diverted traffic and selectively pass a portion thereof to the victim. As explained in paragraph 0301, for example, each guard machine "sieves out the malicious (or excessive) traffic, forwarding to the corresponding victim legitimate traffic at a rate it (the victim) can sustain." As noted in paragraph 0248, "Following filtering and/or at least partial processing of the diverted traffic, some or all of it (e.g., non-malicious packets...) may be directed from the guards to the victim H0)."

Claim 46 recites a network element for use in protecting against an overload condition on a network. The network element includes the following functional components, which may be taken to correspond to elements of the guard machines that are shown in Figures 2 and 3:

(A) An input receives traffic diverted from the network, as represented by the "From Border" arrow in Figure 2. As noted in paragraph 0293, the traffic comprises flows/packets originating from certain IP addresses, i.e., source addresses.

(B) A statistics module performs a statistical analysis of the diverted traffic so as to detect an anomalous pattern of a flow associated with at least one source address. This module corresponds to “statistical engine 16” in Figure 2, which “singles out flows or aggregates of flows (... identified by... the IP addresses...) with irregular/suspicious behavior” (paragraph 0295).

(C) A filter blocks at least a portion of the data packets having the at least one source address, responsively to detection of the anomalous pattern. This filter corresponds to “filter function 12” in Figure 2, which (as explained in paragraph 0293) “blocks packets originating from IP addresses... that were suspected as being a source of malicious traffic...” “The filter-rules are placed in the filter... dynamically by the management of the guard in response to indications received from the statistical engine.”

(D) An output, as represented by the “Back to Border Router” arrow in Figure 2, selectively passes on traffic not blocked by the filter. As noted in paragraph 0248, “Following filtering and/or at least partial processing of the diverted traffic, some or all of it (e.g., non-malicious packets...) may be directed from the guards to the victim H0.”

Claim 49 recites a system for use in protecting against an overload condition on a network. The components of the system, which may be taken to correspond to elements shown in Figure 1, comprise:

(A) One or more network elements (“guards”), shown in Figure 1 as G0-G3. Each guard comprises the following components, which may be taken to correspond to elements shown in Figure 2:

(1) An input for receiving traffic from the network, as represented by the “From Border” arrow in Figure 2.

(2) A filter coupled to the input, which selectively blocks traffic that originated from a source suspected of causing the overload condition. This filter corresponds to “filter function 12” in Figure 2, which (as explained in paragraph 0293) “blocks packets originating from IP addresses... that were suspected as being a source of malicious traffic...”

(3) A statistics module that identifies the traffic statistically indicative of having originated from the source suspected of causing the overload condition.

This module corresponds to “statistical engine 16” in Figure 2, which “singles out flows or aggregates of flows (... identified by... the IP addresses...) with irregular/suspicious behavior” (paragraph 0295).

(4) An output, as represented by the “Back to Border Router” arrow in Figure 2, which selectively passes on to further elements in the network traffic not blocked by the filter. As noted in paragraph 0248, “Following filtering and/or at least partial processing of the diverted traffic, some or all of it (e.g., non-malicious packets...) may be directed from the guards to the victim H0.”

(B) One or more further network elements (“diverters”), shown in Figure 1 as “routers” R0-R8 (paragraph 0247). The routers communicate with the guards and selectively initiate diversion of traffic otherwise destined for a “victim” to at least one of the guards in response to detection of an anomalous traffic condition. This sort of operation is described (with reference to Figure 1) in paragraph 0248, for example: Routers R0-R8 selectively divert traffic destined for a victim H0 to guards G0-G3 when the victim comes under an anomalous traffic condition.

Claim 56 recites a method of responding to an overload condition at a “victim” network element in a set of one or more potential victims on a network. The method includes the following steps:

(A) Traffic destined for the victim is diverted to a guard machine. This sort of operation is described (with reference to Figure 1) in paragraph 0248, for example: Routers R0-R8 selectively divert traffic destined for a victim H0 to guards G0-G3. As noted in paragraph 0293, the traffic comprises flows/packets originating from certain IP addresses, i.e., source addresses.

(B) A statistical analysis of the diverted traffic is performed at the guard machine so as to detect an anomalous pattern of a flow associated with at least one of the source addresses. This step is performed, for example, by “statistical engine 16” in the guard machine that is shown in Figure 2, which “singles out flows or aggregates of flows (... identified by... the IP addresses...) with irregular/suspicious behavior” (paragraph 0295).

(C) Responsively to detecting the anomalous pattern, at least a portion of the packets having the at least one source address that is associated with the anomalous flow pattern are prevented from reaching the victim, while at least some of the packets from other source addresses are passed to the victim. This step is performed, for example, by “filter function 12” of the guard machine shown in Figure 2, which (as explained in paragraph 0293) “blocks packets originating from IP addresses... that were suspected as being a source of malicious traffic...” “The filter-rules are placed in the filter... dynamically by the management of the guard in response to indications received from the statistical engine.” As noted in paragraph 0248, “Following filtering and/or at least partial processing of the diverted traffic, some or all of it (e.g., non-malicious packets...) may be directed from the guards to the victim H0.”

Claim 66 recites a method of responding to an overload condition at a “victim” network element in a set of one or more potential victims on a network. The method includes the following steps:

(A) The victim is coupled to receive traffic from the network via a first port of a network switch. This arrangement is shown in Figure 1, in which victims H0-H4 are connected to routers (i.e., network switches) R6, R4, R0 and R8. Connection of the victims to the routers is described in paragraph 0248.

(B) The network switch is actuated to divert the traffic that is destined for the victim to a second port to which a guard machine is coupled. It can be seen in Figure 1 that guard machine G2 is coupled directly to adjacent router R6 (to which victim H0 is also directly coupled), while other guard machines are remotely coupled to routers R4, R0 and R8 via other routers in the network. Actuation of the routers to divert traffic in this manner to either adjacent or remote guard machines is described in paragraphs 0248-0250.

(C) The diverted traffic is filtered using the guard machine. This step is performed, for example, by “filter function 12” of the guard machine shown in Figure 2.

(D) At least a portion of the filtered traffic is passed selectively from the guard machine to the victim, as described in paragraph 0248: “Following filtering and/or at least partial processing of the diverted traffic, some or all of it (e.g., non-malicious packets...) may be directed from the guards to the victim H0.”

(6) Grounds of Rejection to be Reviewed on Appeal

Claims 1-8, 10, 11, 13-16, 20, 33, 35 and 46-54 and 56-69 were rejected under 35 U.S.C. 103(a) over Jungck (U.S. Patent 6,829,654) in view of Ji et al. (U.S. Patent 6,831,895).

Appellant believes that this ground of rejection should be reversed for all claims.

(7) Argument

I. Independent Claim 1

Appellant respectfully submits that the Examiner erred in maintaining that claim 1 is obvious over Jungck in view of Ji.

Jungck describes apparatus and methods for enhancing network infrastructure using edge servers and edge caches. The edge servers may also be used to detect malicious or otherwise unauthorized data transmissions (abstract). The edge server includes a request interceptor, a request filter and a request transmitter (col. 1, line 62 – col. 2, line 9). The filter prevents any client from directly communicating with any subscribing server (col. 28, lines 40-42) and instead redirects valid content requests to an edge cache (col. 28, lines 21-25). The edge server can also monitor data transmission generated by clients for malicious program code (col. 28, lines 51-57) and can identify the originating client in a DDOS attack (col. 29, lines 3-7). These functions of the edge server are illustrated by Jungck in Fig. 6.

Ji describes a method for relieving congestion by diverting traffic from a congested link to alternative, shortest paths (abstract). These alternative paths take different “hops” through a network to connect a given node to a destination node (col. 4, lines 26-35). Alternative paths of this sort are illustrated by Ji in Figs. 8, 9 and 11-15.

Claim 1 recites a method of responding to an overload condition, in which diversion of traffic by a first set of network elements is initiated in response to an indication of an anomalous traffic condition. The network elements in the first set divert traffic destined for a victim to a second set of network elements, which filter the diverted traffic and selectively pass a portion of the traffic to the victim. This arrangement is advantageous in maximizing network throughput under both normal

conditions and in the presence of an attack, since filtering is applied when and as needed, i.e., in response to an anomalous traffic condition.

In rejecting claim 1, the Examiner acknowledged that Jungck does not teach initiating diversion of traffic due to an indication of an anomalous traffic condition but maintained that this teaching is supplied by Ji. Even if it were conceded, however, that Ji might suggest initiating diversion in response to an anomalous traffic condition, there is no teaching or suggestion in either Jungck or Ji of diverting traffic by a first set of network elements to a second set of network elements, which then filter the traffic, as recited in claim 1.

By contrast to claim 1, as shown by Jungck in Fig. 6, edge servers 602 comprise both the request interceptor 608 and the request filter 606. Although the edge server may redirect certain requests to an edge cache 604, the filtering function is performed not by the edge cache, but rather within the edge server itself. Furthermore, “the request filter 606 pre-filters traffic before receipt by the request interceptor 608” (col. 29, lines 22-23, emphasis added). In other words, not only are the filtering and diversion carried out by the same network element in Jungck, but their order of operation is opposite to that recited in claim 1. Jungck filters in order to decide which traffic to divert, whereas the method of claim 1 diverts traffic in order to filter it.

On the other hand, although Ji describes diverting traffic, he neither teaches nor suggests that the diverted traffic might be filtered or that a portion of the traffic should be passed selectively to its destination. Rather, it is clearly Ji’s intent that all traffic be passed to its destination as expeditiously as possible. Therefore, even if the features of Ji and Jungck were to be combined, the ordered sequence of steps recited in claim 1 – divert and then filter – is outside the range of combinations that a person of ordinary skill could have been motivated to create, except in hindsight from the present patent application.

Thus, the Examiner has failed to make a *prima facie* case of obviousness against claim 1. The cited art neither teaches nor suggests using one set of network elements to divert traffic to a second set of network elements for filtering, nor does it teach or suggest initiating such diversion in response to any sort of traffic condition. Therefore, claim 1 is patentable over the cited art.

II. Independent Claim 46

Appellant respectfully submits that the Examiner erred in maintaining that claim 46 is obvious over Jungck in view of Ji.

Claim 46 recites a network element for use in protecting against an overload condition. The network element comprises an input, a filter for blocking traffic originating from a suspect source, a statistics module, and an output. The statistics module performs a statistical analysis of diverted traffic so as to detect an anomalous pattern of a flow associated with at least one source address. The filter blocks at least a portion of the data packets having such a source address.

In rejecting this claim, the Examiner acknowledged that Jungck does not teach a statistics module that detects an anomalous flow pattern associated with at least one source address, but did not point out the missing teaching in Ji either. Instead, the Examiner simply cited the same passages in Ji (abstract, col. 3, lines 24-59, and col. 4, lines 27-43) as he cited against claim 1. These passages, however, say nothing about any sort of statistics module, let alone a module that performs a statistical analysis of diverted traffic or a module that detects an anomalous pattern of a flow associated with at least one source address, as recited in claim 46. At most, Ji's "IP tuner" determines whether a congested link exists and then adjusts splitting factors associated with the congested link in order to divert traffic to an equal cost path (ECP) (col. 5, lines 34-58). This congested link, however, may be a part of many paths (col. 6, lines 7-9) and is not associated with any particular source address or addresses. Ji neither teaches nor suggests any sort of statistical analysis, let alone the specific type of statistical analysis that is recited in claim 46.

Thus, the Examiner has failed to make a *prima facie* case of obviousness against claim 46. This claim is therefore patentable over the cited art.

III. Independent Claim 49

Appellant respectfully submits that the Examiner erred in maintaining that claim 49 is obvious over Jungck in view of Ji.

Claim 49 recites a system for use in protecting against an overload condition on a network. The system comprises one or more "guards," which comprise an input, a filter for selectively blocking traffic originating from a suspect source, a statistics

module that identifies the traffic statistically indicative of having originated from the suspect source, and an output for passing on traffic not blocked by the filter. One or more “diverters” selectively initiate diversion to the one or more guards of traffic otherwise destined for a victim, responsively to detection of an anomalous traffic condition.

In rejecting this claim, the Examiner asserted that Jungck teaches all the elements of the claim, with the exception of initiating diversion of traffic due to an indication of an anomalous traffic condition. This assertion contradicts the Examiner’s earlier acknowledgment (in regard to claim 46) that Jungck does not teach a statistics module. The earlier acknowledgment was correct: There is nothing at all in Jungck, about statistics, let alone a module that “identifies the traffic statistically indicative of having originated from the source suspected as potentially causing the overload condition,” as recited in claim 49. The passages in Jungck that the Examiner cited as purportedly teaching this sort of module (col. 27, lines 4-46, and col. 29, lines 22-64) relate to a deterministic filtering function: If the packet originated upstream from the edge server, it is considered suspect and is therefore eradicated (col. 29, lines 38-39). Statistical indications play no part at all in the decision.

As explained above in regard to claim 46, Ji likewise fails to teach or suggest the statistics module that is recited in claim 49.

Furthermore, neither Jungck nor Ji teaches or suggests diverting traffic by certain network elements (the “diverters”) to one or more guards, which then filter the traffic, as is recited in claim 49. As explained above in regard to claim 1, not only are filtering and diversion carried out by the same network element in Jungck, but their order of operation is opposite to that recited in claim 49. Jungck filters in order to decide which traffic to divert, whereas the system of claim 49 diverts traffic in order to filter it. Although Ji describes diverting traffic, he neither teaches nor suggests that the diverted traffic might be filtered or that a portion of the traffic should be passed selectively to its destination.

Thus, the Examiner has failed to make a *prima facie* case of obviousness against claim 49. The cited art does not teach or suggest either using one or more network elements to divert traffic to a other network elements for filtering or identifying traffic

statistically as having originated from a suspect source. Therefore, claim 49 is patentable over the cited art.

IV. Independent claim 56

Appellant respectfully submits that the Examiner erred in maintaining that claim 56 is obvious over Jungck in view of Ji.

Claim 56 recites a method for responding to an overload condition, including diverting traffic to a guard machine, performing a statistical analysis of the diverted traffic so as to detect an anomalous pattern of a flow associated with at least one source address, and preventing packets from the source address from reaching the victim. At least a portion of the data packets from other source addresses are passed to the victim.

The Examiner gave no reason for the rejection of claim 56, other than to state that this claim “contain[s] the same language of the claims already discussed above” and is “therefore... rejected under the same rationale” (page 9, fourth paragraph, in the Official Action). As explained above in reference to claims 46 and 49, however, neither Jungck nor Ji teaches any sort of statistical analysis, let alone the specific statistical analysis that is recited in claim 56 for detecting an anomalous pattern of a flow. Furthermore, neither Jungck nor Ji teaches or suggests diverting traffic to a guard machine, which then prevents data packets from a certain source address from reaching a victim.

Thus, the Examiner has failed to make a *prima facie* case of obviousness against claim 56. The Examiner has not met his burden of pointing out specifically how the cited art teaches the claim limitations, and in fact, the cited references do not teach or suggest these limitations. Therefore, claim 56 is patentable over the cited art.

V. Independent Claim 66

Appellant respectfully submits that the Examiner erred in maintaining that claim 66 is obvious over Jungck in view of Ji.

Claim 66 recites a method of responding to an overload condition at a victim network element. The victim is coupled to receive traffic from a network via a first port of a network switch. The network switch is actuated to divert the traffic destined for the victim to a second port, to which a guard machine is coupled. The guard machine filters

the diverted traffic and selectively passes at least a portion of the filtered traffic to the victim.

The Examiner again gave no reason for the rejection of claim 66, other than the blanket statement that this claim “contain[s] the same language of the claims already discussed above” and is “therefore... rejected under the same rationale.” In fact, none of the claims discussed previously by the Examiner in the Official Action makes any reference at all to a network switch, let alone to the use of such a switch to divert traffic to a guard machine in the manner recited in claim 66.

Neither Jungck nor Ji teaches or suggests using a network switch to perform this sort of diversion. As shown in Jungck’s Figs. 6 and 6A, for example, the functions of request filtering, interception, and proxy server are all carried out within the edge server. The edge servers always transmit traffic to the subscribing servers through the same ports. Jungck neither teaches nor suggests that such traffic might be diverted to a guard machine on a different port, for filtering and selective transmission to a victim, as required by claim 66. When Jungck’s edge servers do hand requests off to an edge cache 604, the edge cache satisfies the request itself. The edge cache does not filter and pass the request on to the subscribing server (col. 29, lines 46-50), as would be required by the method of claim 66.

Thus, the Examiner has failed to make a *prima facie* case of obviousness against claim 66. This claim is therefore patentable over the cited art.

There is one other claim that recites the use of a network switch to route traffic destined for a victim to a second port for processing by another network element: dependent claim 55, which the Examiner found to recite allowable subject matter. The Examiner gave no explanation as to why the same subject matter would be allowable in claim 55 but rejected in claim 66. Applicant respectfully submits that this subject matter is patentable in both of claims 55 and 66.

VI. Dependent Claim 3

Appellant respectfully submits that even if claim 1 were conceded to be obvious over the cited art, Jungck and Ji still fail to teach or suggest the added limitations of claim 3.

Claim 3 depends from claim 1, and adds the limitation that the filtering step includes detecting any of a traffic pattern that differs from an expected pattern and a traffic volume that differs from an expected volume in a way that is statistically significant. In rejecting this claim, the Examiner stated that Jungck teaches this limitation in relation to DDoS attacks in col. 28, line 40 – col. 29, line 10. The cited passage, however, has nothing to do with detecting statistically-significant features of traffic patterns or traffic volumes. Rather, Jungck detects attacks in the conventional way, on a packet-by-packet basis, by monitoring packets for malicious program code (col. 28, lines 53-57) or a forged origin address (col. 28, lines 57-62). Jungck neither teaches nor suggests detecting a traffic pattern or volume, and thus says nothing about whether variations of the traffic pattern or volume are statistically significant, as recited in claim 3. Ji likewise makes no suggestion of detecting statistically-significant variations in traffic patterns or volume.

Claim 3 is thus independently patentable over the cited art.

VII. Dependent Claim 7

Appellant respectfully submits that even if claim 1 were conceded to be obvious over the cited art, Jungck and Ji still fail to teach or suggest the added limitations of claim 7.

Claim 7 depends from claim 6, which depends from claim 1, and adds the limitation that the filtering step includes discarding traffic not requiring a selected service from the victim. Examples of this sort of function are described in paragraph 0021 of the present patent application: passing customer orders to the victim while discarding UDP and ICMP packets, or passing mail or IRC packets while discarding other packets.

In rejecting this claim, the Examiner maintained that Jungck teaches the limitations of the claim in col. 29, lines 22-64. The cited passage, however, refers to determining whether the request filter 606 in the edge server 602 should pass certain requests to the edge cache 604 (lines 46-48). If the edge cache is not able to handle the request (lines 50-55) or is not associated with a subscribing server (lines 25-27), the request filter simply passes the traffic through to another destination (lines 27-29 and 53-55). The only condition under which packets might be discarded is if the packets did

not originate from an affiliated POP (lines 31-39). Jungck neither teaches nor suggests discarding traffic that do not require a selected service, as recited in claim 7. Ji does not describe any sort of intentional discard mechanism, let alone a mechanism based on a selected service.

Therefore, claim 7 is independently patentable over the cited art.

VIII. Dependent Claim 8

Appellant respectfully submits that even if claim 7 were conceded to be obvious over the cited art, Jungck and Ji still fail to teach or suggest the added limitations of claim 8. This claim depends from claim 7 and adds that the filtering step includes discarding any of UDP and ICMP traffic. Neither Jungck nor Ji makes the slightest mention of either of these protocols.

In rejecting claim 8, the Examiner asserted that Jungck teaches the limitations of this claim in col. 28, lines 40-46, and col. 31, lines 1-20. The cited passage in col. 28 refers to intercepting traffic having the IP address of a subscribing server, in order to perform value-added services, which may include eradicating forged packets (col. 28, lines 57-62). The cited passage in col. 31 refers to eradicating packets containing unauthorized or malicious program code (lines 15-20) in DNS-based attacks.

Since Jungck makes no mention of either UDP or ICMP, the Examiner's position appears to be that any sort of IP packet discard mechanism necessarily includes filtering out and discarding UDP and ICMP traffic (or evidently, any other protocol that might be carried over IP). The only possible support for this sort of protocol-specific packet discard, however, is impermissible hindsight from the present patent application. Jungck's packet discard criteria are based solely on the source address or malicious program content of the packets. Jungck makes no suggestion that the protocol might be a criterion for packet discard, let alone the specific protocols of UDP and ICMP.

Therefore, claim 8 is independently patentable over the cited art.

IX. Dependent Claim 15

Appellant respectfully submits that even if claim 1 were conceded to be obvious over the cited art, Jungck and Ji still fail to teach or suggest the added limitations of claim 15.

Claim 15 depends from claim 10, which depends from claim 1, and adds the limitation that first and second addresses are associated with the victim. Traffic directed to the first address is discarded if it was received external to an area defined by the points at which the first set of network elements (the traffic diverters) are operated. Traffic directed to the second address, however, is passed to the victim. This “double address” diversion method is described, for example, in paragraphs 0254-0255 *et seq.* in the present patent application.

There is no teaching or suggestion in either Jungck or Ji of assigning two addresses to any sort of network element, let alone using the two addresses in the sort of diversion scheme that is recited in claim 15. The Examiner held that Jungck teaches the limitations of this claim in col. 27, lines 34-51, and col. 28, lines 21-39. The cited passages refer to isolating subscribing servers from network traffic based on source or destination IP addresses (col. 27, lines 38-43, and col. 28, lines 31-34), but do not even hint that multiple addresses might be assigned to these servers for any purpose.

Therefore, claim 15 is independently patentable over the cited art.

X. Dependent Claim 16

Appellant respectfully submits that even if claim 1 were conceded to be obvious over the cited art, Jungck and Ji still fail to teach or suggest the added limitations of claim 16.

Claim 16 depends from claim 10, which depends from claim 1, and adds the limitation that the diverting step includes redirecting traffic using Policy Based Routing. This is a specific type of routing, which is known in the art, based on the incoming interface card of the diverting router, as explained in paragraphs 0267-0268 of the present patent application. Neither Jungck nor Ji makes any mention or suggestion of Policy Based Routing as a possible basis for traffic redirection.

The Examiner held that Jungck teaches the limitations of claim 16 in col. 27, line 13. This passage simply mentions that Jungck’s edge cache 604 may be coupled with routing equipment so as to intercept network traffic. It makes no mention or suggestion of any sort of routing policy.

Therefore, claim 16 is independently patentable over the cited art.

XI. Dependent Claims 20 and 60

Appellant respectfully submits that even if independent claims 1 and 56 were conceded to be obvious over the cited art, Jungck and Ji still fail to teach or suggest the added limitations of claims 20 and 60.

Claim 20 depends from claim 5, which depends from claim 4, which depends from claim 1, and adds the limitation that packets with spoofed source addresses are detected by executing a verification protocol with sources of diverted traffic. Traffic from sources that correctly comply with the verification protocol is passed to the victim. Claim 60 depends from claim 59, which depends from claim 56, and recites that a protocol handshake is initiated between a guard machine and one or more of the source addresses in order to determine spoofed source addresses. In other words, in both of claims 20 and 60, a guard or filtering element interacts with the source of diverted traffic using a certain protocol in order to determine whether the source address is spoofed. Neither Jungck nor Ji teaches or suggests any sort of verification protocol or protocol handshake that might be executed with a traffic source in order to detect spoofed source addresses.

In rejecting claim 20, the Examiner made reference to a verification procedure that is purportedly described by Jungck in col. 28, lines 47 *et seq.* As explained earlier, this passage refers to detecting malicious program code (lines 51-57) and to detecting data packets with implausible origin addresses (lines 57-62). It makes no mention of any sort of protocol that could be used for these purposes, and certainly does not suggest executing a verification protocol or a protocol handshake with a source address that might be spoofed, as recited in claims 20 and 60.

Therefore, claims 20 and 60 are independently patentable over the cited art.

XII. Dependent Claims 35, 54 and 57

Appellant respectfully submits that even if independent claims 1 and 56 were conceded to be obvious over the cited art, Jungck and Ji still fail to teach or suggest the added limitations of claims 35, 54 and 57.

Claim 35, for example, depends from claim 33, which depends from claim 1, and adds the limitation that any of the traffic pattern and volume is determined during a period when the victim is not in an overload condition, for comparison with any of the

traffic pattern and volume in the filtering step (of claim 1) upon detecting the anomalous traffic condition. Neither Jungck nor Ji teaches or suggests this sort of traffic comparison.

The Examiner maintained that Jungck teaches the limitations of claim 35 in col. 29, lines 22-64. The cited passage, however, has nothing to do with determining traffic patterns or volumes, and does not even hint at comparing traffic patterns or volumes under different traffic conditions. Jungck decides which packets to pass to the edge cache or to discard based solely on individual packet characteristics, such as whether the packet originated upstream or downstream (lines 31-38) or contains a request that can be handled by the edge cache (lines 46-48), or contains malicious code (col. 28, lines 51-57). Jungck makes no mention of traffic pattern or volume, and thus cannot possibly be taken to teach or suggest comparing traffic patterns or volumes from different periods as recited in claim 35.

Ji, as noted above, determines whether a congested link exists (col. 5, lines 48-51). There is no teaching or suggestion in Ji, however, of comparing traffic patterns or volumes from different periods as part of a traffic filtering step or for any other purpose.

Claims 54 and 57, which respectively depend from claims 1 and 56, recite that an expected pattern of traffic is learned while the victim is not under attack, and that an anomalous traffic condition or attack is detected when the traffic differs from the expected pattern. As explained above, neither Jungck nor Ji teaches detecting any sort of difference in traffic patterns over time.

Thus, for the reasons explained above, claims 35, 54 and 57 are independently patentable over the cited art.

XIII. Dependent Claim 51

Appellant respectfully submits that even if independent claim 49 were conceded to be obvious over the cited art, Jungck and Ji still fail to teach or suggest the added limitations of claim 51.

Claim 51 depends from claim 49, and adds the limitation that at least one of the guards comprises an ingress filter, coupled to the statistical module, which generates and transmits to another network element rules for blocking traffic on the network. Neither Jungck nor Ji teaches or suggests generation and transmission of any sorts of

rules among network elements, let alone rules for blocking traffic, as recited in claim 51.

In rejecting claim 51, the Examiner cited col. 29, lines 11-64, in Jungck. The cited passage refers to the operation of the request filter 606, which includes “ingress filtering” (line 24), based on whether packets originated upstream or downstream from the edge server (lines 31-38). Jungck does not even hint, however, that the request filter might generate rules or transmit them to other network elements. Furthermore, since Jungck does not describe a statistical module (or any other module that might be considered to perform a statistical function), he cannot possibly be taken to suggest that the ingress filter be coupled to a statistical module, as required by claim 51.

Therefore, claim 51 is independently patentable over the cited art.

XIV. Dependent Claim 53

Appellant respectfully submits that even if claim 1 were conceded to be obvious over the cited art, Jungck and Ji still fail to teach or suggest the added limitations of claim 53. This claim recites diverting all of the traffic destined for the victim upon detecting the anomalous traffic condition.

The Examiner gave no specific grounds for the rejection of claim 53, other than the blanket statement that it contains “the same language of the claims already discussed above.” This statement is incorrect. None of the claims discussed above says anything about diverting all of the traffic that is destined for the victim, as recited in claim 53. This added limitation is neither taught nor suggested by the cited art. On the contrary, Jungck and Ji describe methods and systems in which only certain traffic is diverted, depending on the originating address, packet content, or splitting factors, for example.

Therefore, claim 53 is independently patentable over the cited art.

XV. Dependent Claims 59 and 62

Appellant respectfully submits that even if independent claim 56 were conceded to be obvious over the cited art, Jungck and Ji still fail to teach or suggest the added limitations of claim 59 and 62. These claims depend from claim 56 (directly or indirectly), and add the limitation that data packets with certain source addresses are discarded before performing a statistical analysis of the diverted traffic. In claim 59,

packets with spoofed source addresses are discarded, while in claim 62, it is the packets that have source addresses that are associated with an anomalous traffic flow pattern that are discarded before the statistical analysis.

The Examiner gave no specific grounds for the rejection of claims 59 and 62, other than the same blanket statement that they contain “the same language of the claims already discussed above.” This statement is incorrect. None of the claims discussed above says anything about discarding certain data packets before performing a statistical analysis. This added limitation is neither taught nor suggested by the cited art. Jungck may mention eradicating packets, but does not hint at any sort of statistical analysis. Ji does not suggest discarding packets at any point. There is nothing in the prior art that would have motivated a person of ordinary skill to discard a certain portion of the diverted traffic before performing a statistical analysis of the diverted traffic, as recited in claims 59 and 62.

Therefore, claims 59 and 62 are independently patentable over the cited art.

XVI. Dependent Claim 63

Appellant respectfully submits that even if claim 62 were conceded to be obvious over the cited art, Jungck and Ji still fail to teach or suggest the added limitations of claim 63. This claim depends from claim 62 and recites that after discarding diverted packets that have a source address that is associated with an anomalous flow pattern, the diverted traffic is processed so as to detect and discard data packets with spoofed source addresses. In other words, there are two distinct discard stages before the statistical analysis: (1) packets with source addresses associated with anomalous flow patterns, and (2) packets with spoofed source addresses.

The Examiner again gave no specific grounds for the rejection of claim 63, other than the blanket statement that it contains “the same language of the claims already discussed above.” This statement is incorrect. None of the claims discussed above says anything about carrying out two distinct packet discard stages before performing statistical analysis. There is no disclosure or suggestion of this sort of two-stage discard in the cited art.

Therefore, claim 63 is independently patentable over the cited art.

XVII. Dependent Claim 65

Appellant respectfully submits that even if claim 56 were conceded to be obvious over the cited art, Jungck and Ji still fail to teach or suggest the added limitations of claim 65. This claim depends from claim 56 and adds the limitation that the statistical analysis includes classifying traffic according to types of users that generated the traffic.

The Examiner gave no specific grounds for the rejection of claim 65, other than the blanket statement that it contains “the same language of the claims already discussed above.” This statement is incorrect. None of the claims discussed above says anything about classifying traffic by user type. The only sort of traffic classification in Jungck is by address, and Ji does not teach or suggest any sort of traffic classification.

Therefore, claim 65 is independently patentable over the cited art.

XVIII. Dependent Claim 68

Appellant respectfully submits that even if independent claim 66 were conceded to be obvious over the cited art, Jungck and Ji still fail to teach or suggest the added limitations of claim 68. This claim depends from claim 66 and adds the limitation that after the guard machine filters the diverted traffic, the filtered traffic is passed back from the guard machine to the network switch for transmission to the victim.

The Examiner gave no specific grounds for the rejection of claim 68, other than the blanket statement that it contains “the same language of the claims already discussed above.” This statement is incorrect. None of the claims discussed above says anything about passing filtered traffic back through a network switch that previously diverted the traffic in order to transmit the filtered traffic to its original destination, as recited in claim 68. Ji teaches clearly against such a solution, since the result would be that previously-diverted traffic is redirected back to the congested link that it was supposed to have bypassed.

Therefore, claim 68 is independently patentable over the cited art.

XIX. Remainder of Dependent Claims

Claims 2, 4-6, 10-11, 13-14, 33, 47-48, 50, 52, 58, 61, 64, 67 and 69, which were not expressly discussed above, are patentable over the prior art at least for the reasons set forth above with respect to the claims from which they depend.

Summary

For the foregoing reasons, Appellant submits that the Examiner's rejection of claims 1-8, 10, 11, 13-16, 20, 33, 35 and 46-54 and 56-69 was erroneous. Reversal of his decision is respectfully requested.

Respectfully submitted,

2/11/08
Date

/David J. Powsner/
David J. Powsner
Registration No. 31,868
NUTTER MCCLENNEN & FISH LLP
World Trade Center West
155 Seaport Boulevard
Boston, Massachusetts 02210-2604
(617) 439-2717
(617) 310-9717

APPENDIX A - CLAIMS

The claims involved in the appeal are as follows:

1. A method of responding to an overload condition at a network element ("victim") in a set of one or more potential victims on a network, the method comprising the steps of

A. responsively to an indication of an anomalous traffic condition, initiating diversion of traffic destined for the victim by a first set of one or more network elements external to the set of one or more potential victims to a second set of one or more network elements external to the set of one or more potential victims,

B. the element(s) of the second set filtering traffic diverted in step A ("diverted traffic") and selectively passing a portion thereof to the victim.

2. A method according to claim 1, wherein the initiating step includes effecting a path of traffic that differs from a path that traffic would otherwise take to the victim.

3. A method according to claim 1, wherein

the filtering step includes detecting any of (i) a traffic pattern that differs from an expected pattern and (ii) traffic volume that differ from expected volume, the detecting step includes determining whether any of the traffic pattern and volume varies statistically significantly.

4. A method according to claim 1, wherein the filtering step includes detecting suspected malicious traffic.

5. A method according to claim 4, wherein the detecting step includes detecting packets with spoofed source addresses.

6. A method according to claim 1, wherein the filtering step includes detecting traffic requiring a selected service from the victim.

7. A method according to claim 6, wherein the filtering step includes discarding traffic not requiring the selected service from the victim.

8. A method according to claim 7, wherein the filtering step includes discarding any of UDP and ICMP packet traffic.

10. A method according to claim 1, comprising operating one or more elements of the first set at points on the network around the set of one or more potential victims.

11. A method according to claim 10, comprising operating one or more elements of the second set any of adjacent to or external to one or more elements of the first set.

13. A method according to claim 10, wherein the anomalous traffic condition is indicative of a distributed denial of service (DDoS) attack.

14. A method according to claim 10, comprising selectively activating the one or more elements of the first set by declaring a network address of the victim to be close in network distance to one or more elements of the second set.

15. A method according to claim 10, comprising associating the victim with first and second addresses, and wherein the filtering step includes

discarding traffic received external to an area defined by the points directed to the first address, and

passing to the victim traffic received external to an area directed to the second address.

16. A method according to claim 10, wherein the diverting step includes redirecting traffic using Policy Based Routing.

20. A method according to claim 5, wherein detecting the packets with spoofed source addresses comprises executing a verification protocol with sources of the diverted traffic, and wherein the passing step includes passing to the victim traffic from a source that correctly complies with the verification protocol.

33. A method according to claim 1, wherein the filtering step includes statistically measuring any of a traffic pattern and volume so as to identify any of a source and a type of the overload condition.

35. A method according to claim 33, comprising determining any of the traffic pattern and volume during a period when the victim is not in the overload condition, for comparison with any of the traffic pattern and volume in the filtering step upon detecting the anomalous traffic condition.

46. A network element for use in protecting against an overload condition on a network, the network element comprising:

an input for receiving traffic diverted from the network, the traffic comprising flows of data packets having respective source addresses;

a statistics module that is arranged to perform a statistical analysis of the diverted traffic so as to detect an anomalous pattern of a flow associated with at least one of the source addresses;

a filter, which is operative, responsively to detection of the anomalous pattern, to block at least a portion of the data packets having the at least one of the source addresses; and

an output coupled to the input for selectively passing on to further elements in the network traffic not blocked by the filter.

47. A network element according to claim 46, comprising a termination detection module that at least participates in determining when the overload condition has ended.

48. A network element according to claim 46, comprising an antispoofing element that performs at least one of authenticating and verifying a source of traffic.

49. A system for use in protecting against an overload condition on a network, the system comprising:

one or more network elements ("guards") disposed on the network, each network element having

an input for receiving traffic from the network,

a filter coupled to the input, the filter selectively blocking traffic originating from a source suspected as potentially causing the overload condition,

a statistics module that is coupled to the filter and that identifies the traffic statistically indicative of having originated from the source suspected as potentially causing the overload condition, and

an output coupled to the input for selectively passing on to further elements in the network traffic not blocked by the filter,

one or more further network elements ("diverters") disposed on the network and in communication with the guards, the further network elements selectively initiating, responsively to detection of an anomalous traffic condition, diversion to at least one of the guards traffic otherwise destined for a still further network element ("victim") in a set of one or more potential victims on the network.

50. A system according to claim 49, wherein at least one of the guards comprises a termination detection module that at least participates in determining when the overload condition has ended.

51. A system according to claim 49, wherein at least one of the guards comprises an ingress filter, coupled to the statistics module, that generates and transmits to a further network element on the network rules for blocking traffic on the network.

52. A system according to claim 49, comprising an antispoofing element that any of authenticates and verifies a source of traffic.

53. A method according to claim 1, wherein diverting the traffic comprises diverting all of the traffic destined for the victim upon detecting the anomalous traffic condition.

54. A method according to claim 1, and comprising learning an expected pattern of the traffic while the victim is not under attack, wherein detecting the anomalous traffic condition comprises determining that the traffic differs significantly from the expected pattern.

56. A method of responding to an overload condition at a network element ("victim") in a set of one or more potential victims on a network, the method comprising:

diverting to a guard machine traffic destined for the victim, the traffic comprising flows of data packets having respective source addresses;

performing a statistical analysis of the diverted traffic at the guard machine so as to detect an anomalous pattern of a flow associated with at least one of the source addresses; and

responsively to detecting the anomalous pattern, preventing at least a portion of the data packets having the at least one of the source addresses from reaching the victim while passing to the victim at least some of the data packets from other source addresses.

57. A method according to claim 56, wherein performing the statistical analysis comprises learning an expected traffic pattern of the flows while the victim is not under attack, and detecting an attack by determining that the anomalous pattern differs from the expected traffic pattern.

58. A method according to claim 56, wherein performing the statistical analysis comprises detecting any of a traffic volume, port number distribution, periodicity of requests, packet properties, IP geography, and distribution of packet arrival/size.

59. A method according to claim 56, and comprising processing the diverted traffic so as to detect and discard the data packets that have one or more spoofed source addresses before performing the statistical analysis.

60. A method according to claim 59, wherein processing the diverted traffic comprises initiating a protocol handshake between the guard machine one or more of the source addresses in order to determine that the one or more of the source addresses are spoofed.

61. A method according to claim 56, wherein preventing at least the portion of the data packets comprises filtering out the diverted packets that have the at least one of the source addresses.

62. A method according to claim 61, wherein filtering out the diverted packets comprises discarding the diverted packets that have the at least one of the source addresses before performing the statistical analysis on the diverted traffic that remains after the discarding.

63. A method according to claim 62, and comprising processing the diverted traffic after discarding the diverted packets that have the at least one of the source addresses so as to detect and discard the data packets that have one or more spoofed source addresses before performing the statistical analysis.

64. A method according to claim 56, wherein performing the statistical analysis comprises at least one of analyzing one or more of netflow data, server logs, victim traffic, and traffic volume, and classifying the statistical analysis according to types of users that generated the traffic.

65. A method according to claim 56, wherein performing the statistical analysis comprises classifying the traffic according to types of users that generated it.

66. A method of responding to an overload condition at a network element ("victim") in a set of one or more potential victims on a network, the method comprising:

coupling the victim to receive traffic from the network via a first port of a network switch;

actuating the network switch to divert the traffic destined for the victim to a second port to which a guard machine is coupled;

filtering the diverted traffic using the guard machine; and

selectively passing at least a portion of the filtered traffic from the guard machine to the victim.

67. A method according to claim 66, wherein the network switch comprises a router.

68. A method according to claim 66, wherein selectively passing at least the portion of the filtered traffic comprises passing the filtered traffic from the guard machine to the network switch, for transmission to the victim via the first port.

69. A method according to claim 66, wherein filtering the diverted traffic comprises performing a statistical analysis of the diverted traffic so as to detect an anomalous pattern of a flow associated with at least one source address of the traffic, and responsively to detecting the anomalous pattern, preventing at least a portion of the data packets having the at least one source address.

APPENDIX B – EVIDENCE

None presented.

APPENDIX C – RELATED PROCEEDINGS

None.